

CONSISTENCY OF IMPORTANCE SAMPLING ESTIMATES BASED ON DEPENDENT SAMPLE SETS AND AN APPLICATION TO MODELS WITH FACTORIZING LIKELIHOODS

INGMAR SCHUSTER

Natural Language Processing Group, University of Leipzig

ABSTRACT. In this paper, I proof that Importance Sampling estimates based on dependent sample sets are consistent under certain conditions. This can be used to reduce variance in Bayesian Models with factorizing likelihoods, using sample sets that are much larger than the number of likelihood evaluations, a technique dubbed Sample Inflation. I evaluate Sample Inflation on a toy Gaussian problem and two Mixture Models.

1. INTRODUCTION

This paper broadens the scope of the Importance Sampling estimator by providing proofs that under rather mild conditions, estimates based on dependent sample sets are still consistent. This can be used for variance reduction in certain models classes, namely those that exhibit a factorizing structure in their likelihoods. The paper proceeds as follows. In Section 2, standard Importance Sampling techniques as well as an iterated Importance Sampling scheme, Population Monte Carlo, are reviewed. Section 3 first exemplifies which models qualify as having a factorizing structure and introduces Sample Inflation for these models. Sample Inflation is a technique to artificially blow up the number of samples gained from few likelihood evaluations, thus attaining a much larger set of dependent samples. In Section 4, I proof that Importance Sampling estimates based on dependent samples are consistent, i.e. converge to the integral we are trying to estimate. Section 5 reviews related work from the Population Monte Carlo literature. Finally, Section 6 evaluates Sample Inflation on both a Gaussian toy problem as well as two Dirichlet Mixture Model estimations. In the conclusion, I give directions for future work.

2. IMPORTANCE SAMPLING

The Importance Sampling estimator approximates the mean (alternatively: integral, expected value) H of some function h with respect to some probability density f :

$$\begin{aligned} H &= \int f(x)h(x)dx \\ &= \mathbb{E}_f(h(x)) \end{aligned}$$

E-mail address: schuster@informatik.uni-leipzig.de.

This is achieved by sampling from an auxiliary proposal density q . Say we have acquired a sample set X from q . The Importance Sampling estimator is given by

$$(1) \quad \mathfrak{J}(X) = \frac{1}{|X|} \sum_{x \in X} w(x)h(x)$$

where $w(x) = f(x)/q(x)$ is the weight function (Robert & Casella, 1999). It can be used in case f is not given only proportionally but exactly and is a probability density (i.e. is non-negative and integrates to 1). A necessary condition for Importance Sampling to be unbiased is that $q(x) > 0$ whenever $f(x)h(x) \neq 0$. Its variance is given by $\text{var}_q(\mathfrak{J}(X)) = \sigma_q^2/|X|$ (see Owen, 2013). To ensure finite variance, q has to have heavier tails than f (Robert & Casella, 1999).

However, most times we can only compute f proportionally, as the normalizing constant (also called evidence or marginal likelihood) is unknown. In particular, this is often the case in Bayesian Inference, where the posterior over random variables is given proportionally by the product of prior and likelihood terms. Here, the self-normalized Importance Sampling estimator

$$(2) \quad \mathfrak{J}_n(X) = \frac{1}{w_\Sigma(X)} \sum_{x \in X} w_u(x)h(x)$$

can be used (Robert & Casella, 1999), where $w_u(x) = f(x)/q(x)$ is the unnormalized weight function and $w_\Sigma(X) = \sum_{x \in X} w_u(x)$. A variance estimate is given by $\sum_{x \in X} (w_u(x)/w_\Sigma(X))^2 (h(x) - \mathfrak{J}_n(X))^2$ (Owen, 2013). Both standard and self-normalized Importance Sampling are consistent as a direct consequence of the strong law of large numbers (see Geweke, 1989).

2.1. Population Monte Carlo. I will use the Population Monte Carlo (PMC; Cappé et al., 2004) paradigm in one of the experiments in the evaluation section. As PMC is not well known in the Machine Learning community, I will introduce it here in a very concise way. However, the reader might as well skip this section at first and come back to it before reading section 4. See Cappé et al. (2004) for a thorough introduction to PMC and Douc et al. (2007); Marin et al. (2012); Iacobucci et al. (2010) for newer developments.

The PMC method is based on the observation that proposal distributions for Importance Sampling can depend on previous samples without compromising the validity or (asymptotic) unbiasedness of the estimator (Cappé et al., 2004). PMC works by first generating a population of importance samples (hence the name) from a set of proposal distributions. In each new generation of samples, proposal distributions can be built on previous sample generations. To equalize samples, an Importance Resampling step is introduced whereby each sample in the population is resampled with replacement with a probability proportional to its weight Rubin (1987). A detailed description is given in Algorithm 1. The essential feature of the Algorithm in it is step (a): for each sample and each generation, an individual proposal distribution can be used, the only restriction being that it might not depend on samples from the same generation. In its most naive version (which I will be using), PMC enables choosing the proposal distributions for a new generation such that they are centered on samples from previous generations. Generally speaking, the aim when choosing proposal distributions is to minimize the variance of importance weights - thus avoiding infinite variance of the estimate.

Algorithm 1 Population Monte Carlo Algorithm

Input: initial proposal densities, unnormalized density f , population size p , sample size m

Output: list of m samples

Initialize $S = \text{List}()$

for $t = 1$ **to** T **do**

 Initialize $P = \text{List}()$

 Initialize $W = \text{List}()$

for $i = 1$ **to** p **do**

 (a) select proposal distribution $q_{i,t}$

 (b) generate $x \sim q_{i,t}$ and append it to S

 append weight $f(x)/q_{i,t}(x)$ to W

end for

 normalize W to sum to 1

 resample p values from P with replacement with probability given by the corresponding value in W

 and append samples to S

end for

return S

3. MODELS WITH FACTORIZING LIKELIHOOD TERMS

Assume our generative model has the following structure.

$$\begin{aligned}
 \phi &\sim P(\phi|\alpha_\phi) \\
 \gamma_j &\sim P(\gamma_j|\alpha_\gamma) & \forall j \in [1, \dots, K] \\
 d_i &\sim P(d_i|\phi, \gamma, \alpha_d) & \forall i \in [1, \dots, N]
 \end{aligned}$$

where d_i is the i th data point, there are N data points, each P represents some parameterized family of distributions and $\alpha_\phi, \alpha_\gamma, \alpha_d$ are fixed model parameters. Then the posterior over the latent variables ϕ, γ is given by

$$p(\phi, \gamma|d) \propto P(\phi|\alpha_\phi) \prod_{j=1}^K P(\gamma_j|\alpha_\gamma) \prod_{i=1}^N P(d_i|\phi, \gamma, \alpha_d)$$

where N is the number of data points. Now assume further that the likelihood term for each data point d_i depends exactly on one γ_j ($d_i \not\perp \gamma_j$) and is independent of the other variables in γ ($d_i \perp \gamma_m$ for $m \neq j$). This induces a partition on the data points and allows for further factorization of the likelihood term

$$\prod_{i=1}^N P(d_i|\phi, \gamma, \alpha_d) = \prod_{j=1}^K \prod_{d_i \not\perp \gamma_j} P(d_i|\phi, \gamma_j, \alpha_d)$$

This model structure renders the individual γ_i conditionally independent of each other,

$$(3) \quad \gamma_i \perp \gamma_j | \alpha, \phi, d \text{ for } i \neq j$$

Two model classes satisfying these assumptions are probabilistic matrix factorization (discussed in 3.2) and Dirichlet Mixture Models (discussed in 3.3). First however, I will exemplify an Importance Sampling method, called Sample Inflation, that is applicable whenever the assumptions above hold.

Algorithm 2 Importance Sampling for factorizing models

Input: proposal densities $q_\phi, q_{\gamma_1}, \dots, q_{\gamma_K}$, unnormalized density f , sample size m
Output: tuple (S, W) of m samples and weights
Initialize samples list $S = \text{List}()$
Initialize weights list $W = \text{List}()$
while $\text{len}(S) < m$ **do**
 sample ϕ' according to q_ϕ
 for $j = 1$ **to** K **do**
 sample γ'_j according to q_{γ_j}
 end for
 append $(\phi', \gamma'_1, \dots, \gamma'_K)$ to S
 append $f(\phi', \gamma'_1, \dots, \gamma'_K) / (q_\phi(\phi') \prod_{j=1}^K q_{\gamma_j}(\gamma'_j))$ to W
end while
 $W = W / (\sum_{w \in W} w)$ {for self-normalized IS}

3.1. Sample Inflation for Importance Sampling. A straight forward self-normalized Importance Sampler for models with factorizing likelihoods is given in Algorithm 2. If the density f is actually given in normalized form, the self-normalization at the end ($W = W / (\sum_{w \in W} w)$) can be skipped.

Now consider the following modification of Algorithm 2: instead of only generating one sample γ'_j from q_{γ_j} , generate two samples $\gamma_j^{(1)}, \gamma_j^{(2)}$ and append both $(\phi', \gamma_1^{(1)}, \dots, \gamma_K^{(1)})$ and $(\phi', \gamma_1^{(2)}, \dots, \gamma_K^{(2)})$ to the sample list S (and the accompanying weights to the weight list W). Contrary to first intuition, a set of samples generated this way does not jeopardize consistency, for the corresponding proof see section 5. The likelihood term for the second sample costs as much to compute as the likelihood term for the first sample. As the likelihood term is usually the most expensive part of posterior computation, we get two dependent samples (because the same ϕ' appears in both of them) for the computational price of two independent samples. However, we can take advantage of the likelihood structure to get an overall of 2^K dependent samples. If we sample M initial samples instead, we can construct M^K dependent samples for the price of M likelihood evaluations. This grows very quickly, in fact the growth is polynomial in M and exponential in K .

For ease of illustration, consider $M = 2, K = 2$. The likelihood term for the first sample $(\phi', \gamma_1^{(1)}, \gamma_2^{(1)})$ is

$$\prod_{d_i \not\sim \gamma_1} P(d_i | \phi', \gamma_1^{(1)}, \alpha_d) \prod_{d_i \not\sim \gamma_2} P(d_i | \phi', \gamma_2^{(1)}, \alpha_d)$$

and for $(\phi', \gamma_1^{(2)}, \gamma_2^{(2)})$ we have the likelihood

$$\prod_{d_i \not\sim \gamma_1} P(d_i | \phi', \gamma_1^{(2)}, \alpha_d) \prod_{d_i \not\sim \gamma_2} P(d_i | \phi', \gamma_2^{(2)}, \alpha_d)$$

Reusing the factors computed for the first two samples, we can calculate the likelihoods of two more dependent samples, $(\phi', \gamma_1^{(1)}, \gamma_2^{(2)})$ and $(\phi', \gamma_1^{(2)}, \gamma_2^{(1)})$, almost for free!

This gives rise to the Sample Inflation method, given in Algorithm 3. In the algorithm, I use c as a shorthand ranging over joint samples for the random variables $\gamma_1, \dots, \gamma_K$ and $q_\gamma(c)$ as a shorthand for $q_{\gamma_1}(c_1), \dots, q_{\gamma_K}(c_K)$. A way to think about Sample Inflation is

Algorithm 3 Importance Sampling (Sample Inflation)

Input: proposal densities $q_\phi, q_{\gamma_1}, \dots, q_{\gamma_K}$, unnormalized density f , number of independent proposals for ϕ m , number of likelihood evaluations per independent proposal of ϕ M

Output: tuple (S, W) of $m \cdot M^K$ samples and weights

Initialize $S = \text{List}()$

Initialize $W = \text{List}()$

while $\text{len}(S) < m$ **do**

 sample ϕ' according to q_ϕ

for $j = 1$ **to** K **do**

for $i = 1$ **to** M **do**

 sample $\gamma_j^{(i)}$ according to q_{γ_j}

end for

end for

 compute set C of all M^K possible joint samples from set of tuples $\{(\gamma_1^{(m)}, \dots, \gamma_K^{(m)}) : m \in \{1, \dots, M\}\}$

for $c \in C$ **do**

 append (ϕ', c) to S

 append $f(\phi', c) / (q_\phi(\phi')q_\gamma(c))$ to W {reuse previous likelihood factor computations for f }

end for

end while

$W = W / (\sum_{w \in W} w)$

return (S, W)

that we can use the structure of the problem to get a better approximation of the marginal $f(\phi)$ by averaging over an inflated sample set for γ .

3.2. Matrix Factorization. For illustration purposes I will discuss Factor Analysis. Other examples of Bayesian matrix factorization models include *Gamma Process Nonnegative Matrix Factorization* Hoffman et al. (2010), *Probabilistic Matrix Factorization* Salakhutdinov & Mnih (2007) and Poisson Factorization Gopalan et al. (2013). The Factor Analysis model with k latent factors has the structure

$$d_i = \phi\gamma_i + \epsilon_i$$

Here $d_i \in \mathbb{R}^p$, $\epsilon_i \sim N(0, \Sigma)$ is a residual for some covariance matrix Σ , $\phi \in \mathbb{R}^{p \times k}$ is a factor loading matrix and $\gamma_i \sim N(0, I_k)$ is a vector of latent factors (one for each data point, thus K equals the number of data points). I will not discuss the choice of priors on ϕ and Σ ; for a profound discussion of Factor Analysis, see Dunson (2006). The key observation is that the likelihood of d_i does not depend on γ_j for $j \neq i$ and thus each γ_j is conditionally independent of all the other variables in γ : $\gamma_i \perp \gamma_j | \alpha, \phi, d$ for $i \neq j$. Thus, the assumptions from section 3 are satisfied. Sample inflation in the case of factor analysis works by first sampling proposals ϕ' and possibly Σ' , then sampling M proposals for each γ_i . The likelihood of a single sample $\gamma_i^{(1)}$ then is evaluated as $N(d_i | \phi'\gamma_i^{(1)}, \Sigma')$. Lets say we have two data points d_1, d_2 and two samples for each of the γ_i . The likelihood of the

two joint samples $(\phi', \Sigma', \gamma_1^{(1)}, \gamma_2^{(1)}), (\phi', \Sigma', \gamma_1^{(2)}, \gamma_2^{(2)})$ is

$$\begin{aligned} & N(d_1 | \phi' \gamma_1^{(1)}, \Sigma') N(d_2 | \phi' \gamma_2^{(1)}, \Sigma') \\ & \text{and} \\ & N(d_1 | \phi' \gamma_1^{(2)}, \Sigma') N(d_2 | \phi' \gamma_2^{(2)}, \Sigma'). \end{aligned}$$

From the factors computed for these two samples, we get the likelihood for $(\phi', \Sigma', \gamma_1^{(1)}, \gamma_2^{(2)})$ and $(\phi', \Sigma', \gamma_1^{(2)}, \gamma_2^{(1)})$ using almost no additional computation time. In general we get M^K samples using M likelihood evaluations.

3.3. Dirichlet Mixture Models. In Dirichlet Mixture Models each data point is assumed to be generated by a mixture of K base distributions, where parameters of the base distributions are given by $\gamma_1, \dots, \gamma_K$. A Dirichlet prior is placed on the mixture proportions $\phi^{(1)}$. For each data point d_i a categorical variable $\phi_i^{(2)}$ is drawn, indicating which base distribution it is generated from.¹ The full generative model is

$$\begin{aligned} \gamma_j & \sim G_0(\alpha_\gamma) & \forall j \in \{1, \dots, K\} \\ \phi^{(1)} & \sim \text{Dir}(\alpha_\phi) \\ \phi_i^{(2)} & \sim \text{Cat}(\phi^{(1)}) \\ d_i & \sim P(d_i | \gamma_{\phi_i^{(2)}}) \end{aligned}$$

where G_0 is a prior on the parameters of the K base distributions, $\phi^{(1)} \in \mathbb{R}_+^K$, $\phi_i^{(2)} \in \{1, \dots, K\}$ and $P(\cdot | \gamma_{\phi_i^{(2)}})$ is the base distribution with index $\phi_i^{(2)}$ (each base distribution could also have some global parameter α_d , which I drop for notational clarity). Again, observe that d_i does not depend on γ_j for $j \neq \phi_i^{(2)}$ and the assumptions from section 3 hold because $\gamma_i \perp \gamma_j | \alpha_\gamma, \alpha_\phi, \phi^{(1)}, \phi^{(2)}, d$ for $i \neq j$. To apply Sample Inflation to Dirichlet Mixture Models, one would first sample a proposal $\phi'^{(1)}$ and $\phi_i'^{(2)}$ for each i , then M proposals for each γ_j , and recombine these to get M^K dependent samples.

4. RELATED WORK

To the best of my knowledge, a recombination of Importance Samples as suggested in this paper has not been proposed before.

Generally speaking, variance reduction is an important topic in Importance Sampling and its descendant Population Monte Carlo. I will concentrate on the PMC case here. In the original paper by Cappé et al. (2004), the approach used for variance reduction is to keep several markov transition kernels which generate new samples centered on previous ones with a different variance for each kernel. Those kernels which exhibit smaller weight variance are then used more often. Mixture-PMC (M-PMC; Cappé et al., 2008) goes one step further in that it fits a Gaussian or Multivariate t mixture model to the samples from previous generations, generating new samples from this approximation of the posterior. D-Kernel PMC by Douc et al. (2007) fits a D-Kernel Mixture and can be shown to converge to the optimum D-Kernel Mixture.

¹The notation differs from the usual notation in DP Mixture Models. However, I valued consistency with section 3 higher than consistency with the rest of the literature.

5. IMPORTANCE SAMPLING ESTIMATORS BASED ON DEPENDENT SAMPLES

In the literature, the sample set used for the Importance sampling estimator is often assumed to consist only of independently identically distributed (iid) samples. However, one potentially interesting (and as we will see practically very relevant) case is when samples are guaranteed to come from the proposal density q but are not required to be independent. I will first introduce some assumptions and notation for this section.

Definition 1. Let X_1, \dots, X_k with fixed k be (multi-)sets of samples from some density (for claims about the Importance Sampling estimator, from the proposal density q). The samples in each X_i are assumed to be iid but the samples in the (multi-)set $X_\cup = \bigcup_i X_i$ are not necessarily independent. Furthermore, let $X_\cup^{(m)} = \bigcup_{i=1}^k X_i^{(m)}$ be a sequence of sample sets with fixed k , $m = |X_\cup^{(m)}|$ and $|X_i^{(m)}| \xrightarrow{m \rightarrow \infty} \infty$. The samples in each $X_i^{(m)}$ are assumed to be iid for any m and i , but the samples in $X_\cup^{(m)}$ might be dependent.

Now as a first step towards proving consistency of Importance Sampling estimates based on dependent sample sets, we note that the normed error of any convex combination of estimates based on iid sample sets cannot increase compared to the same convex combination of normed errors of individual estimates.

Theorem 1. Let $\hat{\mathcal{T}}$ be any estimator of the true quantity H . Then the normed error of a convex combination of estimates $\sum_{i=1}^k \lambda_i \hat{\mathcal{T}}(X_i)$ cannot exceed the convex combination of normed errors:

$$\sum_{i=1}^k \lambda_i \|\hat{\mathcal{T}}(X_i) - H\| \geq \left\| \sum_{i=1}^k \lambda_i \hat{\mathcal{T}}(X_i) - H \right\| \geq 0$$

for any norm $\|\cdot\|$ and $\sum_{i=1}^k \lambda_i = 1, \forall i : \lambda_i \geq 0$. In particular, this implies the squared error of the convex combination of estimators cannot exceed the convex combination of squared errors.

Proof. We have

$$\begin{aligned} & \sum_{i=1}^k \lambda_i \|\hat{\mathcal{T}}(X_i) - H\| \\ \geq & \left\| \sum_{i=1}^k \lambda_i (\hat{\mathcal{T}}(X_i) - H) \right\| \\ = & \left\| \left(\sum_{i=1}^k \lambda_i \hat{\mathcal{T}}(X_i) \right) - \left(\sum_{i=1}^k \lambda_i H \right) \right\| \\ = & \left\| \sum_{i=1}^k \lambda_i \hat{\mathcal{T}}(X_i) - H \right\| \\ \geq & 0 \end{aligned}$$

where first inequality follows either from subadditivity and absolute homogeneity of norms or from Jensens inequality and the fact that norms are convex. The second inequality follows from the positivity property of norms. \square

Now I will specialize Theorem 1 to the case of the (normalized) Importance Sampling estimator. Recall that we are trying to estimate the integral $H = \int f(x)h(x)dx$.

Theorem 2. *The normed error of the estimate $\mathfrak{J}(X_{\cup})$ cannot exceed the cardinality weighted average of normed errors:*

$$\sum_{i=1}^k \frac{|X_i|}{|X_{\cup}|} \|\mathfrak{J}(X_i) - H\| \geq \|\mathfrak{J}(X_{\cup}) - H\| \geq 0$$

where $|\cdot|$ signifies the cardinality of a set.

Furthermore, the normed error of the estimate $\mathfrak{J}_n(X_{\cup})$ cannot exceed the importance weighted average of normed errors:

$$\sum_{i=1}^k \frac{w_{\Sigma}(X_i)}{w_{\Sigma}(X_{\cup})} \|\mathfrak{J}_n(X_i) - H\| \geq \|\mathfrak{J}_n(X_{\cup}) - H\| \geq 0.$$

Proof. For the case of the unnormalized estimator \mathfrak{J} , if we choose $\lambda_i = |X_i|/|X_{\cup}|$ and show $\sum_{i=1}^k \frac{|X_i|}{|X_{\cup}|} \mathfrak{J}(X_i) = \mathfrak{J}(X_{\cup})$, the claim follows from Theorem 1. Using the definition of the estimator (1) we have

$$\begin{aligned} \sum_{i=1}^k \frac{|X_i|}{|X_{\cup}|} \mathfrak{J}(X_i) &\stackrel{(1)}{=} \sum_{i=1}^k \frac{1}{|X_{\cup}|} \sum_{x \in X_i} w(x) h(x) \\ &= \frac{1}{|X_{\cup}|} \sum_{x \in X_{\cup}} w(x) h(x) \\ &\stackrel{(1)}{=} \mathfrak{J}(X_{\cup}) \end{aligned}$$

and thus the first claim holds.

For the case of the self-normalized estimator \mathfrak{J}_n , if we choose $\lambda_i = w_{\Sigma}(X_i)/w_{\Sigma}(X_{\cup})$ and show $\sum_{i=1}^k \frac{w_{\Sigma}(X_i)}{w_{\Sigma}(X_{\cup})} \mathfrak{J}_n(X_i) = \mathfrak{J}_n(X_{\cup})$, the claim follows from Theorem 1. Using the definition of the estimator (2) we have

$$\begin{aligned} \sum_{i=1}^k \frac{w_{\Sigma}(X_i)}{w_{\Sigma}(X_{\cup})} \mathfrak{J}_n(X_i) &\stackrel{(2)}{=} \sum_{i=1}^k \frac{1}{w_{\Sigma}(X_{\cup})} \sum_{x \in X_i} w_u(x) h(x) \\ &= \frac{1}{w_{\Sigma}(X_{\cup})} \sum_{x \in X_{\cup}} w_u(x) h(x) \\ &\stackrel{(2)}{=} \mathfrak{J}_n(X_{\cup}) \end{aligned}$$

and thus the second claim holds. \square

To get an intuition for the meaning of Theorem 2 for the case of the unnormalized estimator, recall that by using Algorithm 3, we can get M^K iid sample sets. Each of these is of size m , so the convex combination amounts to a simple average. Thus, we can only do better on average by using the samples from all sets as compared to the samples from only one set. This seems particularly fortunate after realizing that there is no reason to prefer one of the sample sets over one of the others (all of them are sampled iid from q).

Now if the normed error cannot increase when using dependent sample sets for estimation, we might expect that the estimate converges in probability to the true integral H . In other words, we might expect that the sequence of estimates is consistent. This is indeed the case as stated by the following theorem.

Theorem 3. *Let $\hat{\mathfrak{J}}$ be the (self-normalized) Importance Sampling estimator. The sequence of estimates $\hat{\mathfrak{J}}(X_{\cup}^{(m)})$ for $m \rightarrow \infty$ is consistent, i.e.*

$$\lim_{m \rightarrow \infty} \mathbb{P}(\|\hat{\mathfrak{J}}(X_{\cup}^{(m)}) - H\| \geq \epsilon) = 0$$

for all $\epsilon > 0$.

Proof. By standard results for all i and any $\epsilon > 0$:

$$\lim_{|X_i^{(m)}| \rightarrow \infty} \mathbb{P}(\|\widehat{\mathfrak{J}}(X_i^{(m)}) - H\| \geq \epsilon) = 0$$

Now assume that This implies that for any convex combination

$$\lim_{m \rightarrow \infty} \mathbb{P}\left(\sum_{i=1}^k \lambda_i^{(m)} \|\widehat{\mathfrak{J}}(X_i^{(m)}) - H\| \geq \epsilon\right) = 0$$

using $\sum_i \lambda_i^{(m)} \epsilon = \epsilon$. Now if $\widehat{\mathfrak{J}} = \mathfrak{J}$ choose $\lambda_i^{(m)} = |X_i^{(m)}|/|X_{\cup}^{(m)}|$, if $\widehat{\mathfrak{J}} = \mathfrak{J}_n$ choose $\lambda_i^{(m)} = w_{\Sigma}(X_i^{(m)})/w_{\Sigma}(X_{\cup}^{(m)})$ and apply Theorem 2 to get

$$\lim_{m \rightarrow \infty} \mathbb{P}(\|\widehat{\mathfrak{J}}(X_{\cup}^{(m)}) - H\| \geq \epsilon) = 0$$

□

An important detail of Theorem 3 is that the number of sets k is fixed as $m \rightarrow \infty$.

One of the major reasons for choosing Importance Sampling over other simulation techniques is that it enables approximating model evidence (also called marginal likelihood or the normalizing constant of f). This is based on the identity

$$F = \int f(x) dx = \int \frac{f(x)}{q(x)} q(x) dx = \mathbb{E}_q \left(\frac{f(x)}{q(x)} \right)$$

which yields the unbiased and consistent estimator

$$(4) \quad \mathfrak{Z}(X) = \frac{1}{|X|} \sum_{x \in X} f(x)/q(x)$$

Evidence estimates based on sample sets that contain dependent samples will stay consistent as follows from Theorem 3 by setting $h(x) = 1$.

6. EVALUATION

In this section, I will evaluate Sample Inflation for two cases. As a very simple measure, we will look into the performance of Sample Inflation when computing the (known) expectation of a two dimensional Gaussian distribution with diagonal covariance matrix. As a more involved case we will consider the estimation of two Dirichlet Mixture Models.

6.1. Expectation of a multivariate Gaussian. For the two experiments in this subsection, 20,000 samples were drawn from the respective two dimensional proposal distribution. These were unchanged for standard Importance Sampling estimation. For Sample Inflation, the sample set was partitioned into sets of 100 samples, which were inflated and concatenated into 2,000,000 dependent samples.

As a first evaluation case I chose a multivariate normal, $f = N(0, 2I)$, as the target distribution. The log evidence (log normalizing constant) was artificially set to -1000 . The proposal distribution was a multivariate t -Distribution with the same mean and covariance matrix and 20 degrees of freedom, $q = T(0, 2I, 20)$. Squared bias, mean squared error (MSE) and variance of the estimates for the expectation of the target distribution as well as the evidence are given in log-log-plots in Figure 1. For estimation of the targets expectation, the MSE, which subsumes variance and squared bias, clearly shows that using Sample Inflation is preferable to standard Importance Sampling. The picture is less clear cut for evidence approximation, but Sample Inflation does not seem hurt performance strongly.

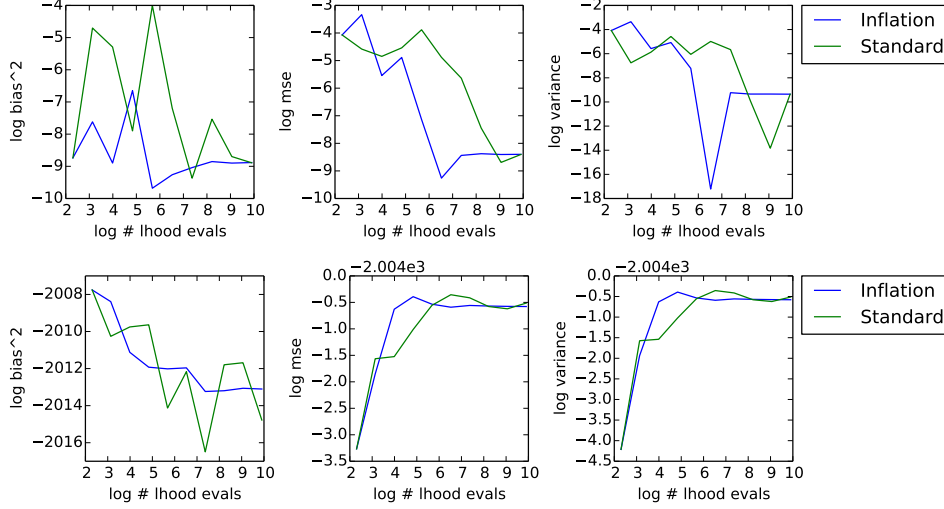


FIGURE 1. Performance of Sample Inflation as compared to standard self-normalized Importance Sampling from the same proposal distribution. The top plots give squared bias, mean squared error and variance for estimation of the expectation of the target distribution. The bottom plots plot the same measures for estimation of the evidence. The proposal distribution was centered on the true expectation of the target distribution.

The second experiment used the same target, $f = N(0, 2I)$ with a log evidence of -1000 . This time however, the proposal distribution was not centered on the target, but on $(5, 5)^T$, $q = T((5, 5)^T, 2I, 20)$. The mean squared error evaluation does not favor Sample Inflation for estimation of the targets expectation this time, though Sample Inflation gives more stable estimates (Figure 2). The major contribution to MSE here comes from the bias, which is caused by the fact that our proposal distribution is not centered on the target. For evidence approximation, Sample Inflation hurts performance slightly, but bear in mind that the differences to standard Importance Sampling are small when transformed back from log space.

6.2. Estimation of DMMs. In this evaluation, I use a Population Monte Carlo approach to estimate two Dirichlet Mixture Models (DMMs) for synthetic data sets comprised of 100 data points. For both experiments, 2000 samples were drawn for standard Importance Sampling. For Sample Inflation, after sampling $\phi^{(1)}$ and $\phi^{(2)}$, two dependent samples were drawn for the parameters of the two component distributions (thus $M = 2, K = 2$). I used less overall samples for Sample Inflation than for standard Importance Sampling, so as to keep the number of likelihood evaluations exactly equal.

In the first case, the synthetic data was generated from a mixture of two one dimensional Gaussians with different means and unit variance. The DMM used two Gaussian components with fixed unit variance. Thus, only the means of the components had to be estimated. I put an $N(0, 1)$ prior on the component means. I used Gaussian Markov kernels to generate proposals based on samples from earlier generations of the PMC algorithm. Sample

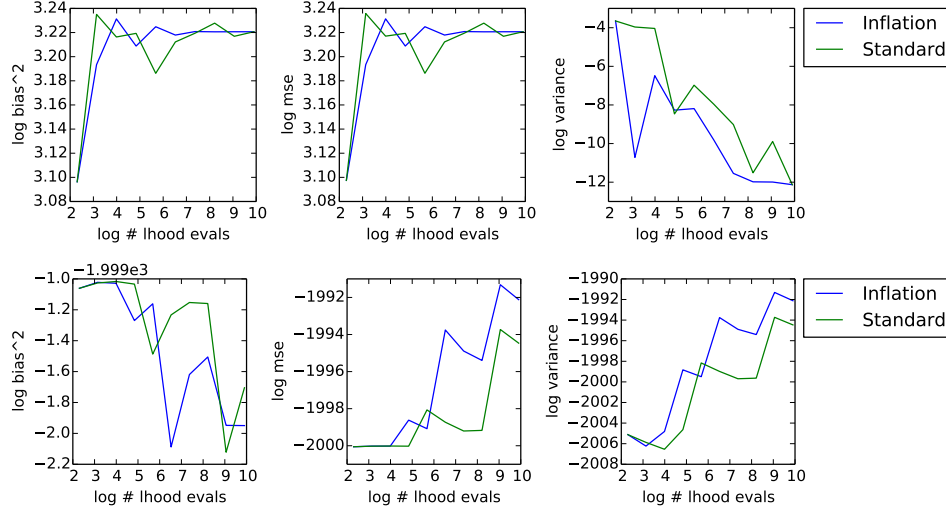


FIGURE 2. Performance of Sample Inflation as compared to standard self-normalized Importance Sampling from the same proposal distribution. The top plots give squared bias, mean squared error and variance for estimation of the expectation of the target distribution. The bottom plots plot the same measures for estimation of the evidence. The proposal distribution was *not* centered on the true expectation of the target distribution.

Inflation attained regions of high likelihood more quickly and exhibited lower variance than standard Importance Sampling (Figure 3).




FIGURE 3. Gaussian Mixture Model estimation. Solid lines mark Sample Inflation, dashed lines standard Importance Sampling. The true means of the synthetic data are dotted. Using Sample Inflation, high likelihood regions are reached more quickly. This is reflected in the estimated means, which are closer to the true means of the data for Sample Inflation.

In the second case, the synthetic data was generated from a mixture of two one dimensional T-Distributions with different means, unit variance, and 30 degrees of freedom. The DMM used two T components. Based on a sample S from a previous generation, I used T-distributed Markov kernels to generate proposals for the mean centered on value of the mean parameter in S . Equivalently, I used Inverse-Wishart Distributions centered on the covariance matrix in S and Gamma distributions centered on the degrees of freedom in S . A Student- t $T(0, 1, 1)$ prior was placed on the component means, an Inverse Wishart $IW(\sigma^2 = 5, \text{df} = 1)$ prior on the covariance and a $\text{Gamma}(1, 1)$ prior (shape and scale parametrization) on the degrees of freedom. Here, Sample Inflation is much better in achieving high likelihoods more quickly, though the estimates exhibit higher variance (Figure 4). The estimates of the means are not close to the true means of the synthetic data, which probably stems from the fact that the Mixture Model is very flexible as we also estimate the covariance matrix and degrees of freedom. Also, the $IW(\sigma^2 = 5, \text{df} = 1)$ prior on the covariance and the $\text{Gamma}(1, 1)$ prior on degrees of freedom are very broad and compensate easily for the rather narrow $T(0, 1, 1)$ prior on the component means.




FIGURE 4. T Mixture Model estimation. Solid lines mark Sample Inflation, dashed lines standard Importance Sampling. The true means of the synthetic data are dotted. Using Sample Inflation, high likelihood regions are reached much more quickly.

7. CONCLUSION

The contributions of this paper were twofold. First, I proved that Importance Sampling estimates based on dependent sample sets are consistent under mild conditions. To the best of my knowledge, this has not been proved before or if it has, the mainstream literature does not reflect this. Second, I apply this to models with factorizing likelihoods, resulting in Sample Inflation, a technique to generate many dependent samples from few likelihood evaluations. The evaluation in section 6 showed that Sample Inflation can reduce variance and help to attain high likelihood regions more quickly in a Population Monte Carlo setting. Future work will have to derive variance estimates for Sample Inflation and, as a consequence, measures of Effective Sample Size and perplexity Robert & Casella (2010). This will hopefully lead to a better understanding of when Sample Inflation can help and under which conditions it hurts performance.

ACKNOWLEDGMENTS

I thank Ilja Klebanov for proof reading and very valuable suggestions with regard to notation and readability. Patrick Jähnichen and Shirin Riazzy checked the proofs and provided helpful discussions, as did Christoph Teichmann. Christian Robert hinted at certain improvements in readability.

REFERENCES

- Cappé, O, Guillin, a, Marin, J. M., and Robert, C. P. Population Monte Carlo. *Journal of Computational and Graphical Statistics*, 13(4):907–929, December 2004. ISSN 1061-8600. doi: 10.1198/106186004X12803. URL <http://www.tandfonline.com/doi/abs/10.1198/106186004X12803>.
- Cappé, Olivier, Douc, Randal, Guillin, Arnaud, Marin, Jean-Michel, and Robert, Christian P. Adaptive importance sampling in general mixture classes. *Statistics and Computing*, 18(4):447–459, April 2008. ISSN 0960-3174. doi: 10.1007/s11222-008-9059-x. URL <http://link.springer.com/10.1007/s11222-008-9059-x>.
- Douc, R., Guillin, a., Marin, J. M., and Robert, C. P. Convergence of adaptive mixtures of importance sampling schemes. *Annals of Statistics*, 35(1):420–448, 2007. ISSN 00905364. doi: 10.1214/009053606000001154.
- Dunson, David B. Efficient Bayesian model averaging in factor analysis. Technical report, National Institute of Environmental Health Sciences, 2006. URL <ftp://ftp.stat.duke.edu/pub/WorkingPapers/06-03.pdf>.
- Geweke, John. Bayesian inference in econometric models using Monte Carlo integration. *Econometrica*, 57(6):1317–1339, 1989. ISSN 0012-9682. URL <http://www.jstor.org/stable/pdfplus/1913710.pdf>.
- Gopalan, Prem, Hofman, Jake M, and Blei, David M. Scalable Recommendation with Poisson Factorization. *arXiv preprint*, pp. 1–10, 2013. URL <http://arxiv.org/pdf/1311.1704>.
- Hoffman, Matthew D, Blei, David M, and Cook, Perry R. Bayesian Nonparametric Matrix Factorization for Recorded Music. *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, pp. 439–446, 2010. doi: 10.1.1.168.2855. URL http://machinelearning.wustl.edu/mlpapers/paper_files/icml2010_HoffmanBC10.pdf<http://www.icml2010.org/papers/523.pdf>.
- Iacobucci, Alessandra, Marin, Jean-Michel, and Robert, Christian. On variance stabilisation in Population Monte Carlo by double Rao-Blackwellisation. *Computational Statistics & Data Analysis*, 54(3):698–710, March 2010. ISSN 01679473. doi: 10.1016/j.csda.2008.09.020. URL <http://linkinghub.elsevier.com/retrieve/pii/S0167947308004489>.
- Marin, Jean-michel, Pudlo, Pierre, and Sedki, Mohammed. Consistency of the Adaptive Multiple Importance Sampling. pp. 1–31, 2012. URL <http://arxiv.org/abs/1211.2548>.
- Owen, Art B. *Monte Carlo theory, methods and examples*. Book Draft, 2013.
- Robert, Christian and Casella, George. *Introduction Monte Carlo Methods with R*. Springer, 2010. ISBN 9781441915757.
- Robert, Christian P. and Casella, George. *Monte Carlo statistical methods*, 1999. URL <http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Monte+Carlo+Statistical+Methods#0>.

- Rubin, Donald B. *Multiple Imputation for Nonresponse in Surveys*. Wiley Publishing, 1987. ISBN ISBN: 047108705X ; 9780471087052; Series ISSN: 0271-6232; LCCN: 86-28935. doi: 10.1002/9780470316696.
- Salakhutdinov, Ruslan and Mnih, Andriy. Probabilistic Matrix Factorization. In *Neural Information Processing Systems*, pp. 1–8, 2007. ISBN 9781605582054. doi: 10.1145/1390156.1390267. URL <http://discovery.ucl.ac.uk/63248/http://discovery.ucl.ac.uk/63251/>.